Design and Analysis of the On-chip Power Delivery Network for Energy Efficient Designs

Sept 21, 2012

Kyle Craig

Advisors: Benton H. Calhoun, John Lach

Robust Low Power VLSI

Motivation

- Many applications impose energy constraints
- High Performance energy constraint:
  - Thermal
- Low Performance energy constraint:
  - Battery lifetime

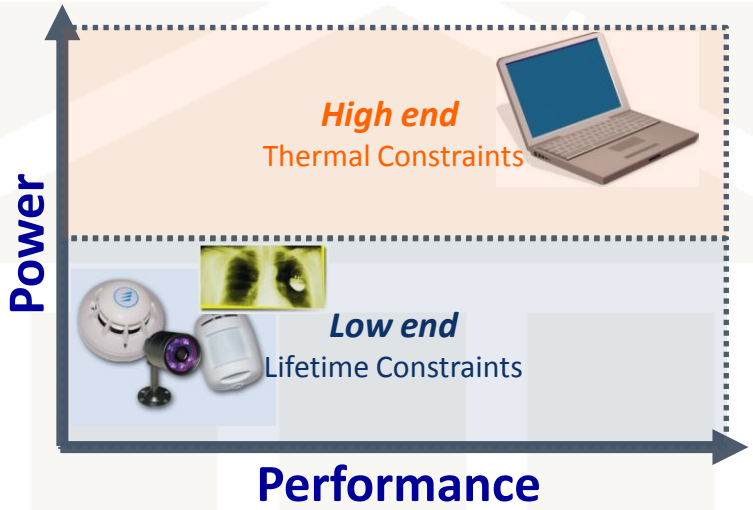

Energy efficiency is one of the largest focuses in digital integrated circuit design

Power delivery network

- Simplified model
  - Voltage generation
  - Printed circuit board (PCB)
  - Package
  - Chip
- Resistances, inductances, capacitances

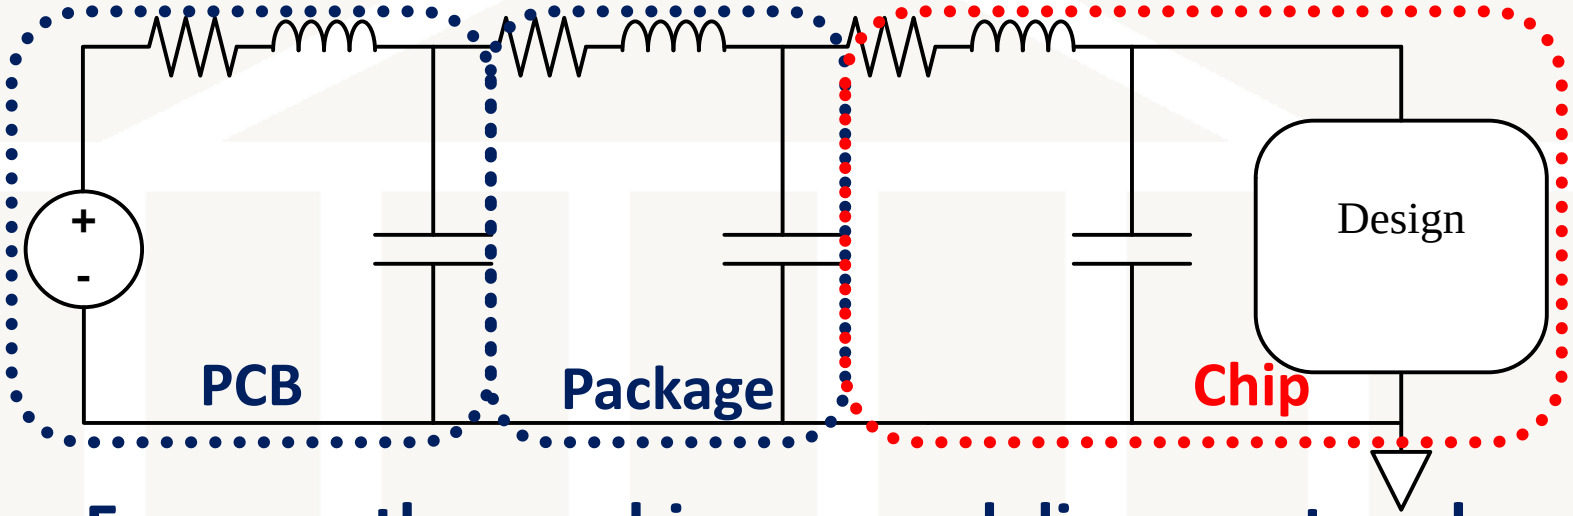

Focus on the on-chip power delivery network

Prior Art

- Dynamic voltage and frequency scaling
  - Conventionally off-chip power delivery network optimization

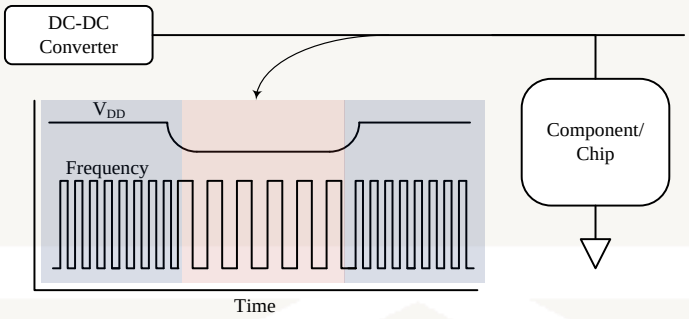

$$E_{op}(V_{DD}) = C_{eff}(V_{DD}) * V_{DD}^2 + V_{DD} * I_L(V_{DD}) * t_{op}$$

- Power gating
  - On-chip power delivery network optimization

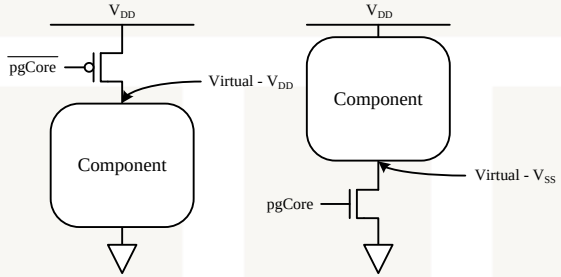

- Subthreshold operation
  - Operate below device threshold voltage (V_T)
  - Quadratic energy reduction, exponential performance decrease
  - Mode change i.e., infrequent voltage scaling

Design challenges power delivery network

- Voltage scaling
  - Slow off-chip DC-DC converters
- V_{DD} granularity
  - Shared common V_{DD} across multiple cores
- IR drop
  - Voltage drop across power gates
- di/dt noise
  - Rush current

Limit the opportunity for V_{DD} scaling, reducing energy efficiency

Research goals

1. Improve voltage scaling architectures/power delivery network
   - Move V_{DD} scaling from off-chip to on-chip
     - Use power switches
   - Improve V_{DD} granularity
2. Enable subthreshold operation
3. Power delivery network noise analysis
   - Evaluate proposed optimizations impact
   - Propose noise mitigation techniques
4. Scripted design space exploration tool

Outline

I. Motivation
II. Background
III. Research goals
IV. Voltage scaling architectures
   I. Motivation
   II. Panoptic dynamic voltage scaling (PDVS)
   III. A programmable resistive power grid
V. Enabling subthreshold operation
VI. Power delivery network noise analysis
VII. Scripted design space exploration tool
VIII. Schedule
IX. Publications

Voltage scaling focus

- Motivation: Enable energy efficient operation
  - Move V_{DD} scaling on-chip
  - Improve V_{DD} granularity
- Implement PDVS [1] in a processor
  - Compare to other power delivery network optimizations
- A programmable resistive power grid
  - Leverage IR drop for energy savings
  - Create model for design

Panoptic dynamic voltage scaling

- Modify on-chip power delivery network
  - Discrete set of voltage rails
  - PMOS power switches for voltage scaling
- Fine V_{DD} granularity
  - Component level

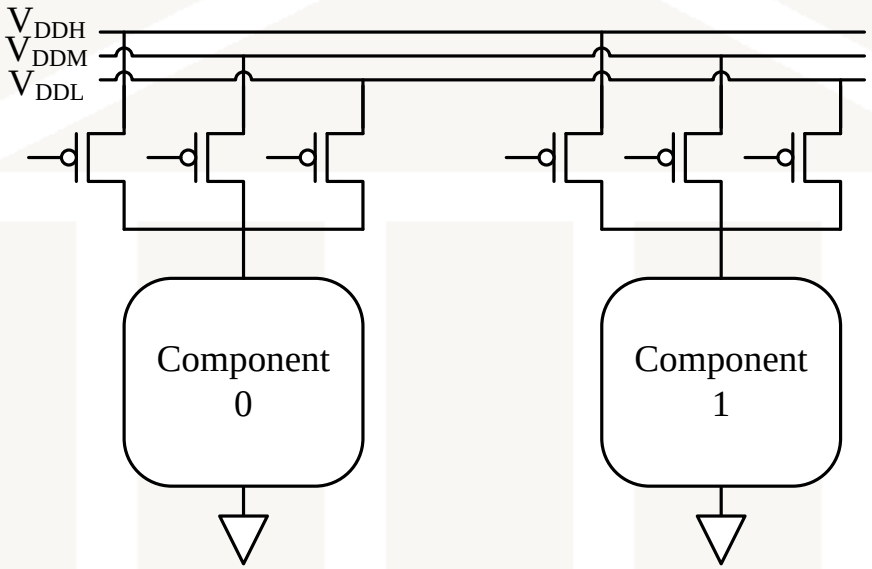

Approach

- 32 bit data flow processor
  - Adder/Multipliers attached to PMOS headers
- Fair comparison to single-V_{DD}, multi-V_{DD}

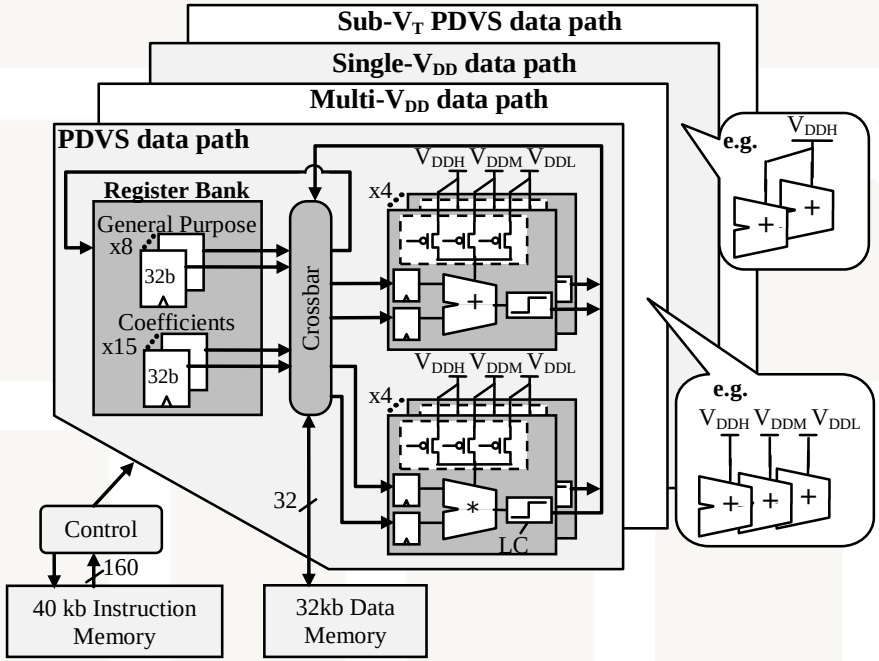

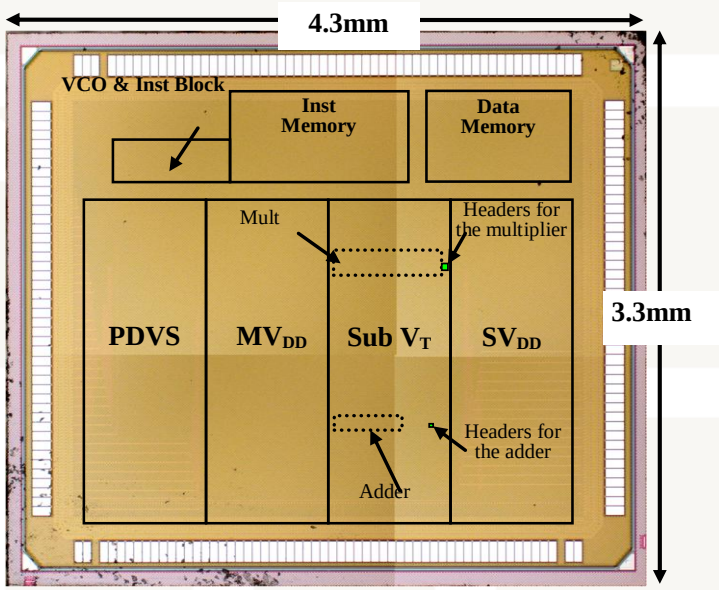

Energy savings up to 50% and 46% over single-V_{DD} & multi-V_{DD}

Source: [2]

Programmable resistive power grid

- Monolithic header (/footer) broken into partitions (w^k_n)
  - Independent gate control → control resistance
  - W^k total width
- Fewer headers on → increased IR drop

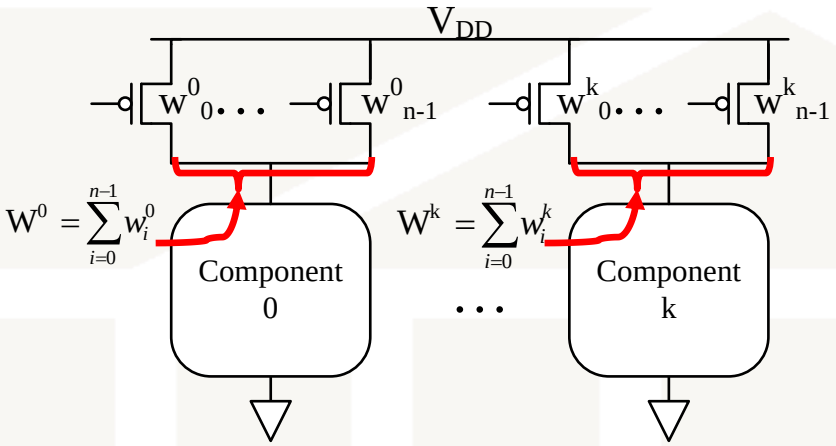

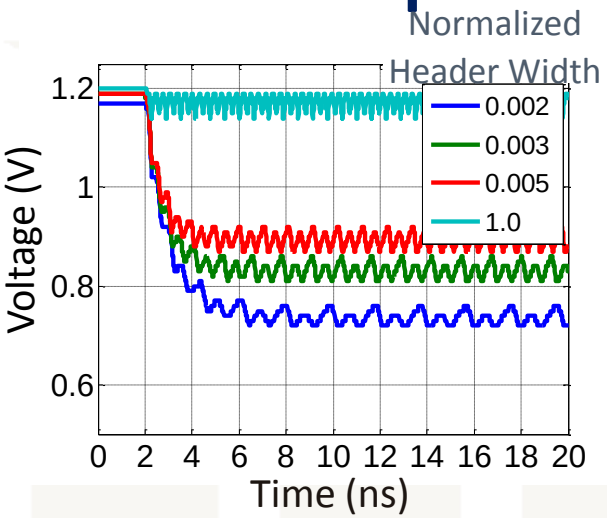

- Modified energy equation

$$E_{op}(V_{DD}, VV_{DD}) = C_{eff}(VV_{DD}) * V_{DD} * VV_{DD} + V_{DD} * I_L(VV_{DD}) * t_{op}$$

Source: [3]

Multi-core model

- AMD Bulldozer core
- Route Level Macro (RLM)'s
  - Building blocks of core
  - Time dependent current source and capacitance
- Apache Redhawk
  - Commercial power integrity tool
- Simplified RLC model

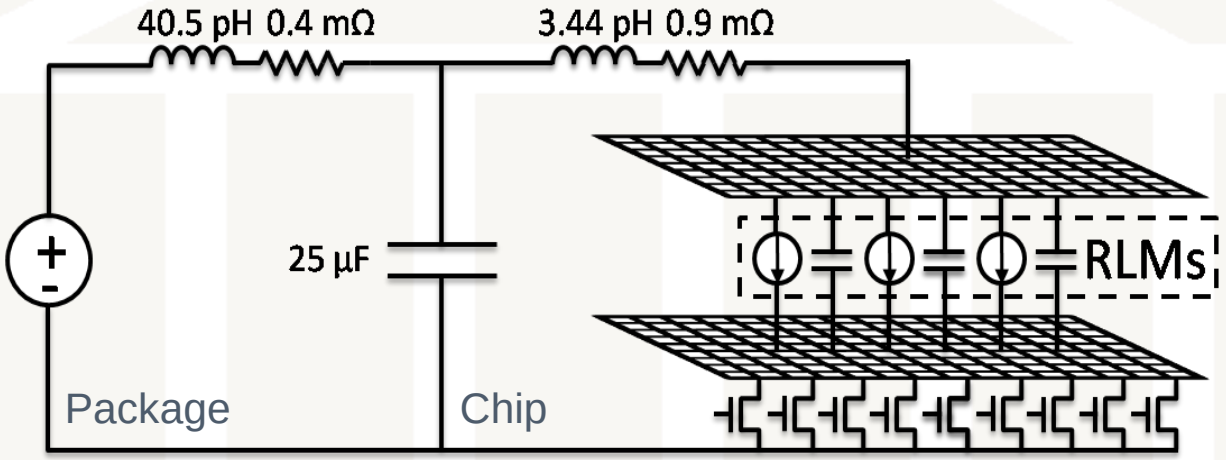

Source: [3]

Model results

- Double-precision General Matric Multiply (DGEMM) benchmark
  - 25ns timing window
  - All RLMs superimposed

Small change in footer partition, significant change in Virtual-V_{SS}

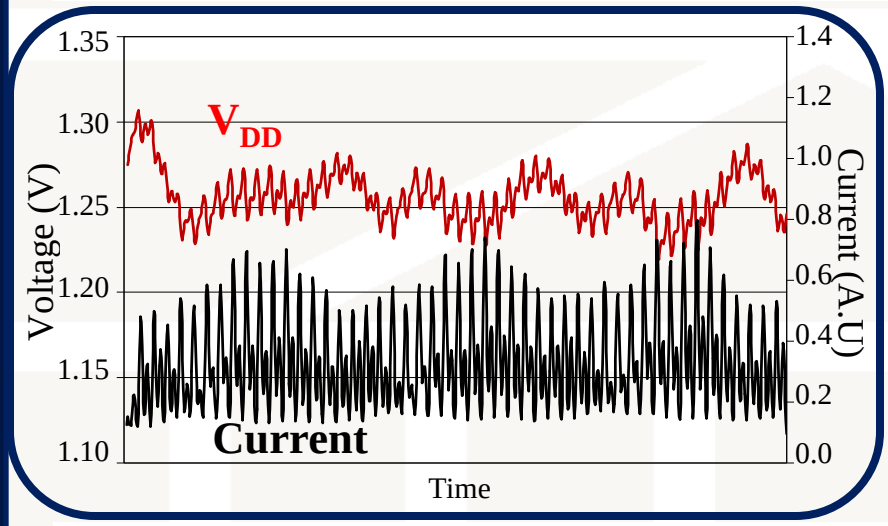

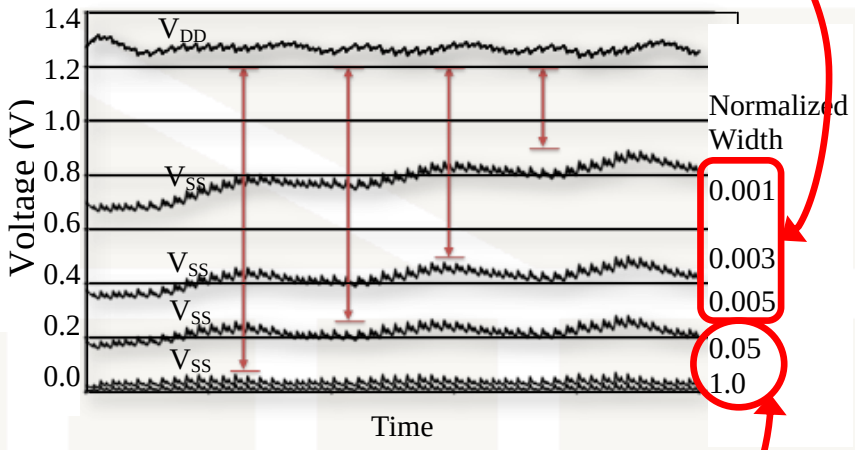

Large change in footer partition. Insignificant change in Virtual-V_{SS}

V_{DD} response shows our model is properly working

Source: [3]

Proposed contributions

- First processor implementing PDVS
  - Energy savings up to 50% and 46% over single-V_{DD} & multi-V_{DD}
- Programmable resistive power grid
  - Leverage IR drop for energy efficiency
- Programmable resistive power grid model
  - Design, implementation and verification

Outline

I. Motivation
II. Background
III. Research goals
IV. Voltage scalable architectures
V. Enabling subthreshold operation
   I. Motivation
   II. Architecture changes in PDVS
   III. On-chip power delivery network optimizations
VI. Power delivery network noise analysis
VII. Scripted design space exploration tool
VIII. Schedule
IX. Publications

Enabling subthreshold focus

- Motivation: Enable energy efficient operation through subthreshold
  - Infrequent mode change
- Architecture modifications
  - In PDVS
- Power delivery network modifications
  - NMOS vs. PMOS for subthreshold header
  - Transmission gate vs. NMOS or PMOS for flexible designs

PDVS: Enabling subthreshold

- Added PMOS to power delivery network of:
  - Register bank
  - Crossbar
- Added bypass for component level converters
  - Level converters not designed for subthreshold operation
- Optimized level converter for subthreshold [4]
  - Converts from 0.25 up to 1.0V
  - Added bypass from data path to memories

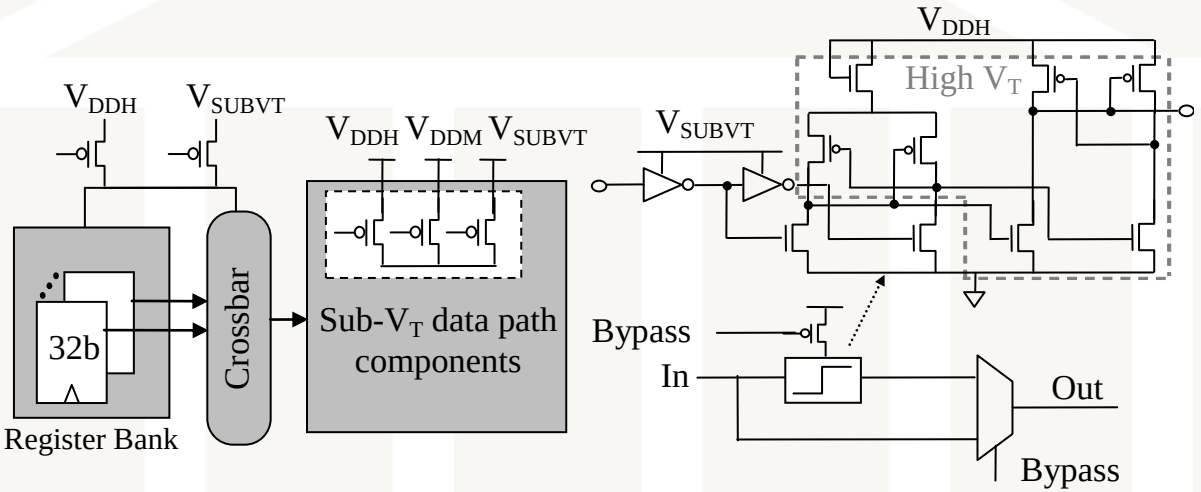

Source: [2]

PMOS vs. NMOS

- PMOS transistor
- During subthreshold operation

$|V_{GS}| = V_{SUBVT}$

- NMOS transistor
- **Gate at V_{DDH}**
- During subthreshold operation

$|V_{GS}| = V_{DDH} - V_{SUBVT}$

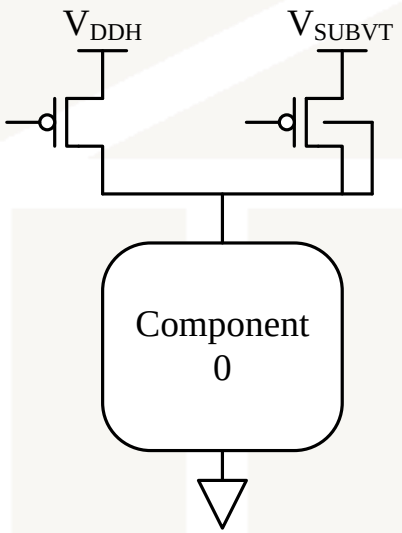

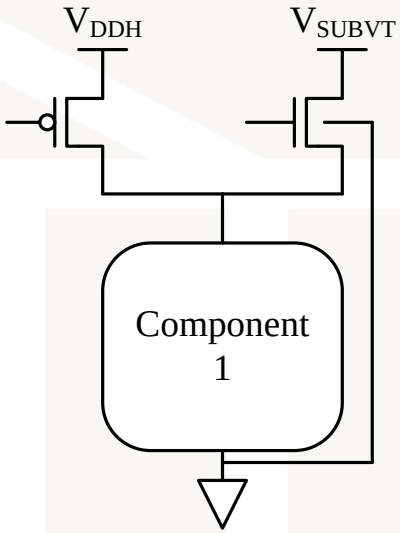

NMOS provides more current

Approach

- Commercial 130nm bulk CMOS process
- PMOS and NMOS design consisted of:
  - Varied PMOS/NMOS width
  - Ten 27-stage ring oscillators (RO) in parallel
    - Each RO can be enabled independently
    - 10 enabled => activity of 1, 1 enabled => activity of 0.1
  - V_{DDH} => 1.2V, V_{SUBVT} => 0.3V

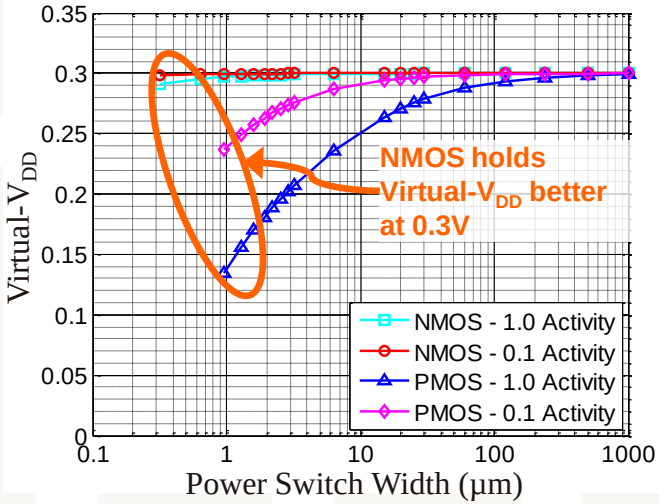

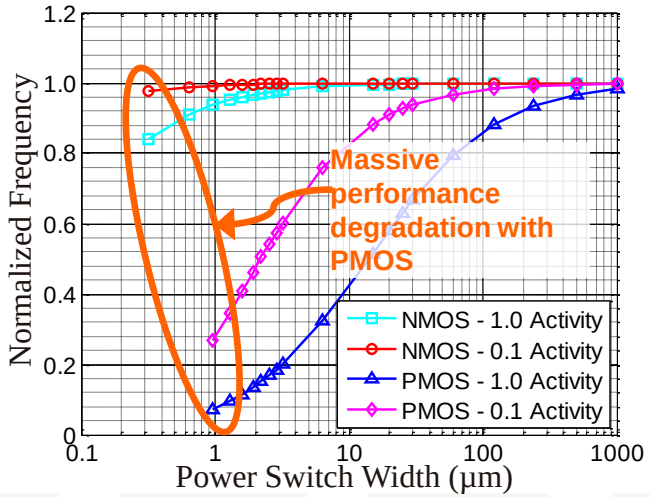

NMOS *280X smaller* than PMOS for a 10% performance degradation

Source: [4]

Transmission gate for flexibility

- Provide flexibility for designs that have a scalable V_{DDL} voltage rail
  - When $V_{SUBVT} > V_T$, PMOS dominate device
  - When $V_{SUBVT} < V_T$, NMOS dominate device
- Same setup previously described
  - NMOS & PMOS width constant
  - Varied V_{SUBVT} voltage

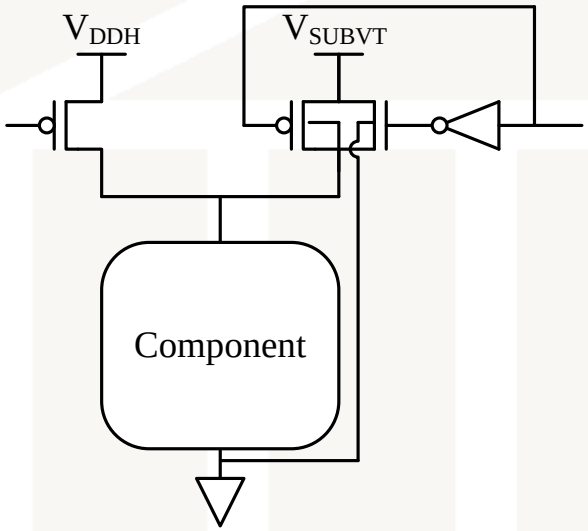

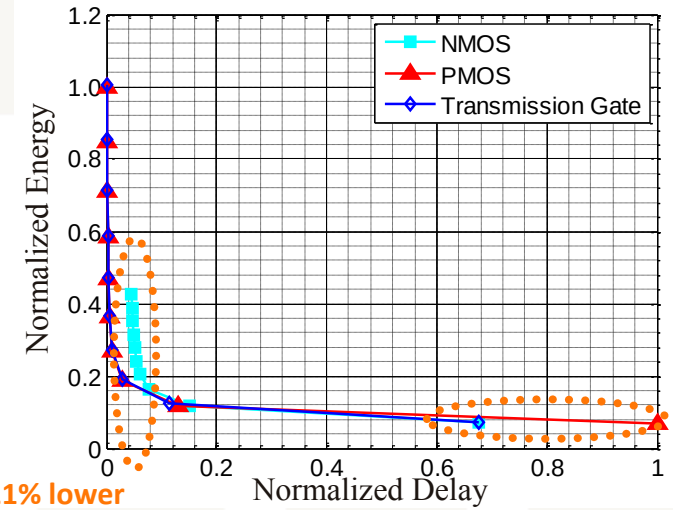

~33% faster for the same energy

21% lower energy for the same delay

Source: [5]

Proposed contributions

- Methodology for adapting architectures for sub-threshold operation
  - PDVS used as example
- Proposed use of NMOS as subthreshold header
  - Gate at V_{DDH}
- Proposed use of transmission gate for flexible designs.
  - V_{SUBVT} can be super or subthreshold
- Comparison between PMOS, NMOS, and transmission gate headers across V_{DD}

Outline

I. Motivation
II. Background
III. Research goals
IV. Voltage scalable architectures
V. Enabling subthreshold operation
VI. Power delivery network noise analysis
   I. Motivation
   II. Panoptic dynamic voltage scaling
   III. Field programmable core array (FPCA)
VII. Scripted design space exploration tool
VIII. Schedule
IX. Publications

Power delivery noise analysis focus

- Motivation: Analyze and evaluate impact of our proposed techniques on the entire power delivery network.

- Synthesis and place & route approach
- Analysis of physical impacts of PDVS
  - V_{DD} switching, i.e., V_{DDL} to V_{DDH}
  - Power grid design for varied V_{DD}s
- Analysis of physical impacts of field programmable core array (FPCA)
  - Impacts of reconfigurable architecture
- Use Encounter Power System (EPS)
  - Commercial power integrity tool

Synthesis and place & route approach

- Synthesis: behavioral RTL→ structural RTL
  - Cadence RTL compiler (RC)
- Place and route: structural RTL → physical layout
  - Cadence Encounter
- Limitations:
  - Native support for power gating, V_{DD} domains, clock gating

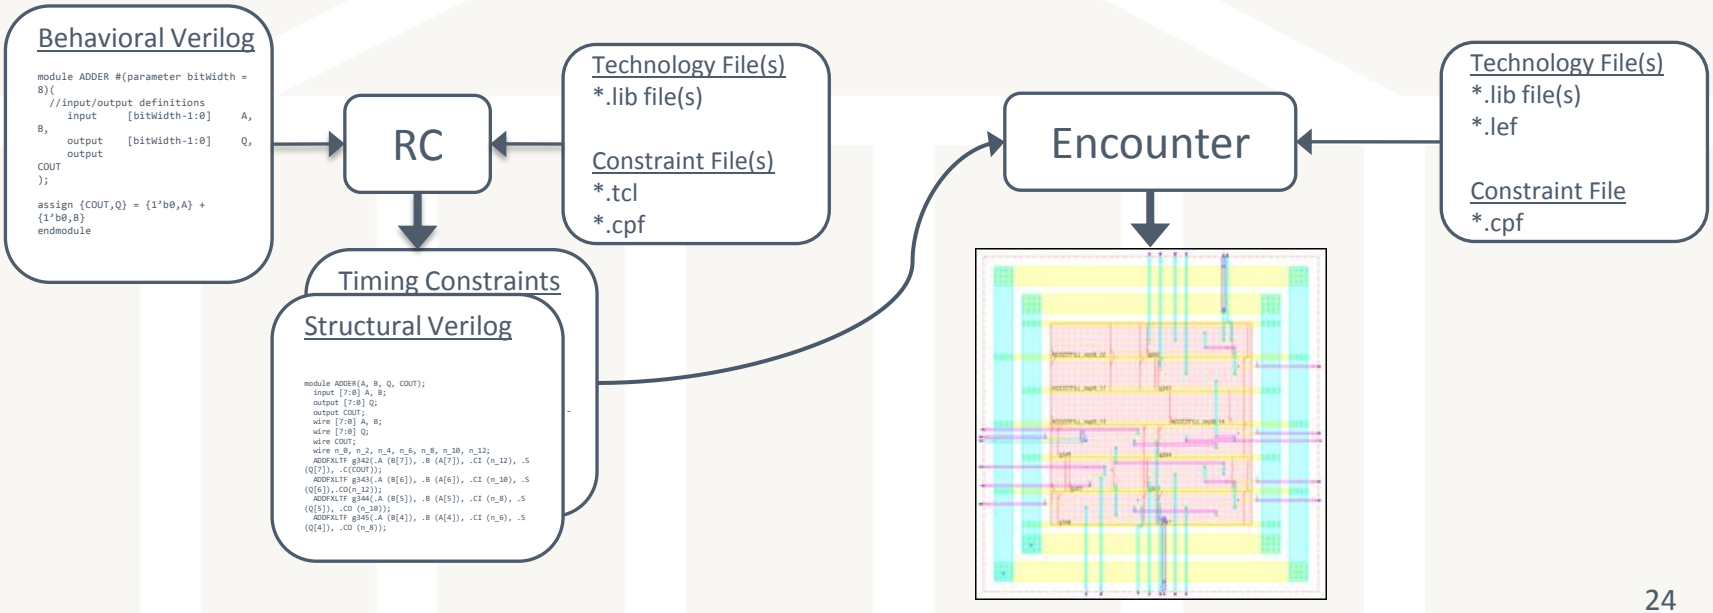

PDVS design approach

- Modified tool flow
  - ‘careful manipulation’
  - Scripted
- Example: PIC processor
  - ALU → 3 PMOS headers – V_{DDH}, V_{DDM}, V_{DDL}
  - Everything else → V_{DDH}

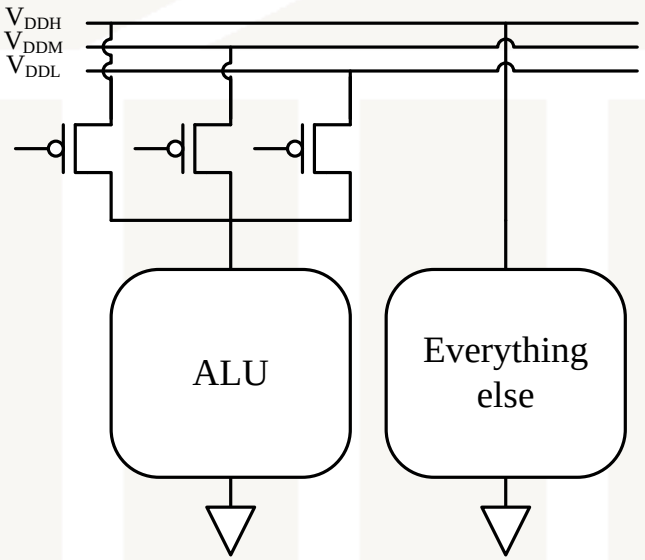

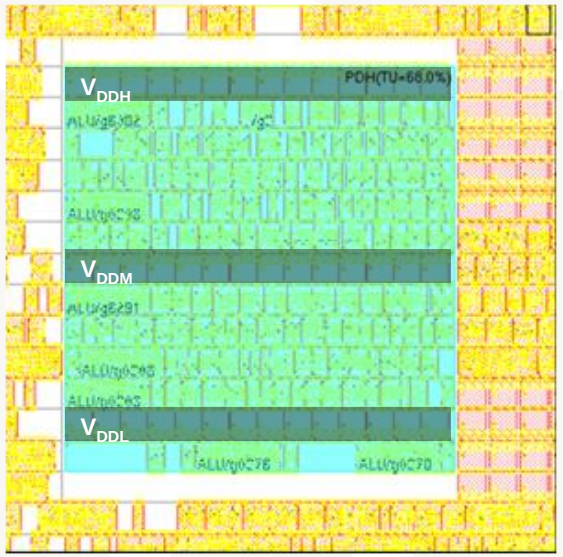

FPCA overview

- Reconfigurable architecture
- OpenRISC
  - Split into front end (FE) & processing element (PE)
  - Reconfigurable interconnect
- Variable width SIMD
- Switch between SIMD and MIMD

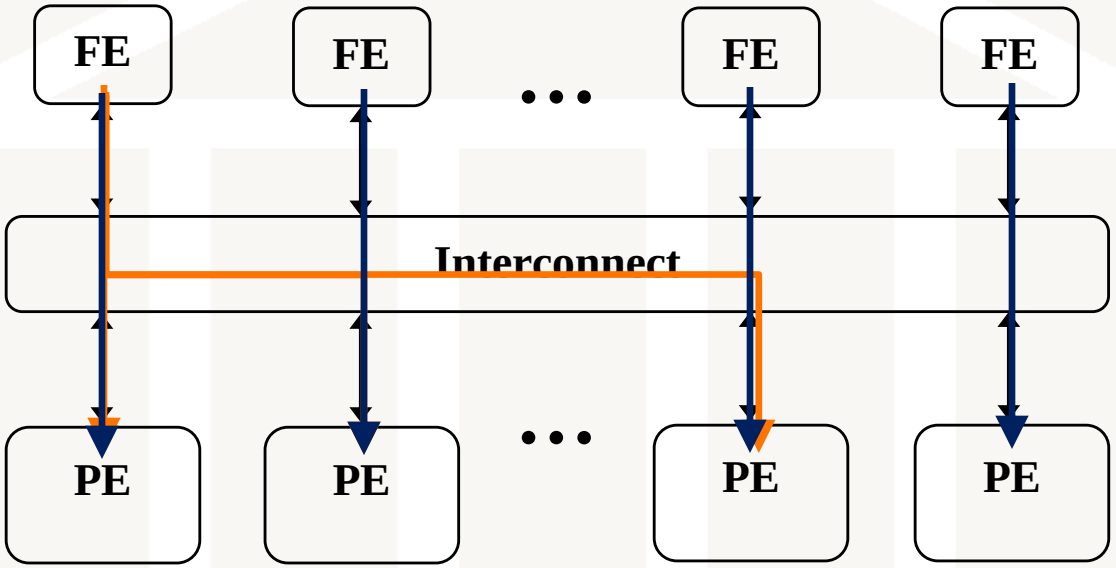

FPCA approach

- Synthesis and place and route flow
- Design size prohibitive
- Need tiled, scalable, hierarchical approach
  - Use interface logic module (ILM)
  - Allows us to maintain timing, speed up place & route

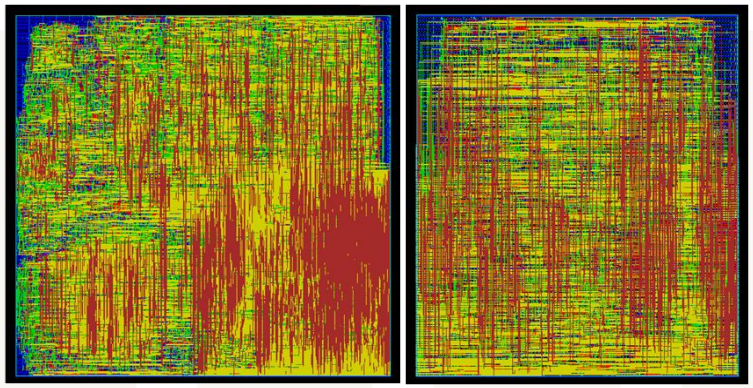

PE tile

FE tile

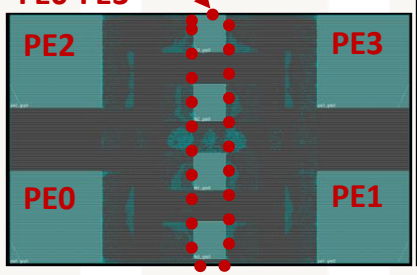

4FE4PE

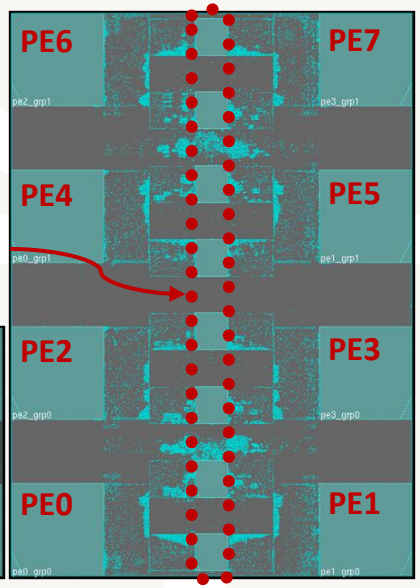

8FE8PE

Encounter Power System approach

- Power delivery network integrity analysis tool
- Static and dynamic analysis
  - IR drop across power gates in PDVS & FPCA
- Rush current analysis (di/dt)
  - PDVS – switching from V_{DDL} to V_{DDH}
  - FPCA – varying SIMD width
- Simplified RLC package model
- Design knobs
  - V_{DD} voltage, metal allocation, power switch size & distribution
- Example: power gated PIC processor
  - IR map, IR histogram

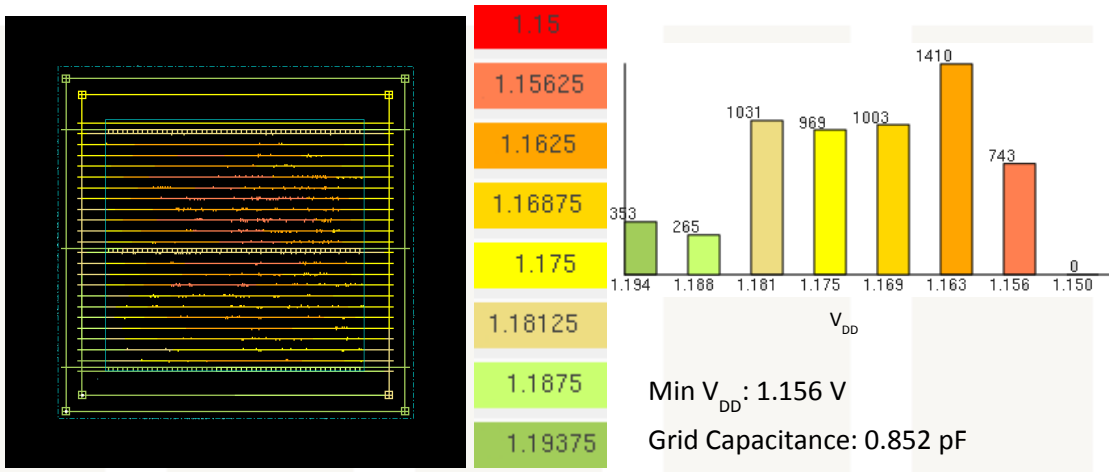

Proposed contributions

- PDVS synthesis and place & route flow
- Analysis and methodology for characterizing power delivery network properties associated with implementing PDVS
- Power grid design methodology for varied voltages
- Methodology for characterizing the power delivery network for reconfigurable architectures

Outline

I. Motivation
II. Background
III. Research goals
IV. Voltage scalable architectures
V. Enabling subthreshold operation
VI. Power delivery network noise analysis
VII. Scripted design space exploration tool
VIII. Schedule
IX. Publications

Tool flow focus

- **Motivation:** Need a tool flow that allows for rapid design space exploration

- Current design flow
  - GUI based
  - Energy efficient techniques natively supported
    - Power gating, clock gating, V_{DD} domains
- Proposed design flow
  - Scripted
  - Easy to use

GUI based place and route

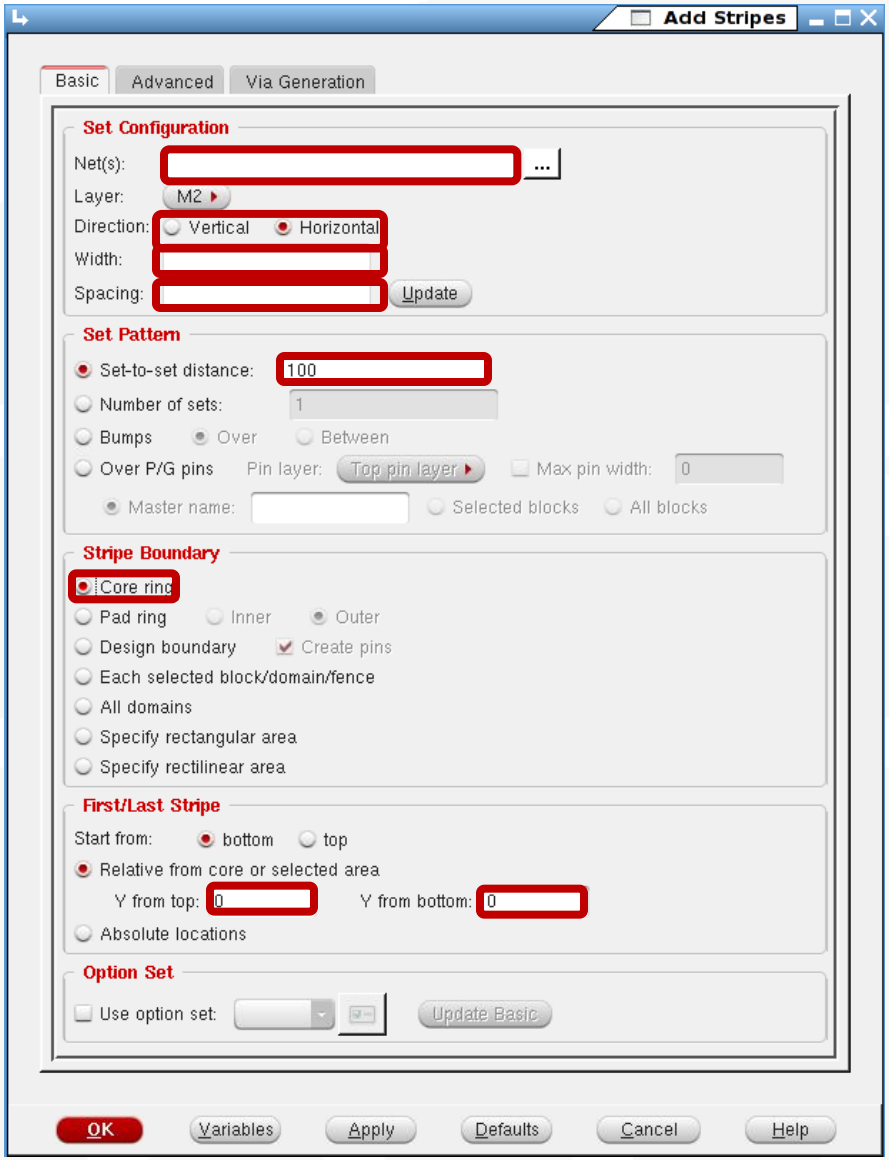

- Creating power grid
- Need to set many fields

For every metal layer

GUI based place and route

- Steps of place & route
  - Tool set up
    - Import: RTL, constraints, library files, LEF, etc..
  - Floor planning
  - Power planning
  - Standard cell placement
  - Clock tree synthesis
  - Routing
  - Verification

Cumbersome, error prone, time consuming

Proposed design tool

- Scripted
  - TCL language, perl
  - Easy to use, command line based
    - ‘launchrc.pl –powergate 0 –clockgate 1’
- Support for proposed energy efficient techniques
  - PDVS
  - Programmable resistive power grid
  - NMOS header
  - FPCA

Proposed contributions

- Scripted design tool enabling energy efficient design space exploration
- Include standard energy efficient techniques
  - Power gating
  - Clock gating
  - V_{DD} domains
- Include proposed energy efficient techniques
  - PDVS
  - Programmable resistive power grid
  - NMOS header
  - FPCA

Schedule

Subject	#	Task description	Status	Related publications
PDVS	1	Design space exploration	Completed	
	2	Simulation	Completed	
	3	Layout	Completed	
	4	Chip testing	Completed	KAC1, KAC2, KAC3
	5	Subthreshold testing	Sep-2012	KAC7
	6	Synthesis, place and route, noise analysis	Apr-2013	KAC9, KAC10
Resistive power grid	7	Design space exploration	Completed	
	8	Redhawk Modeling	Completed	KAC4
	9	Layout - from AMD internship	Completed	
	10	Layout - from BSN2	Nov-2012	
	11	Testing AMD chip	? Still waiting on silicon	
	12	Testing BSNrev2	Jun-2013	KAC8, KAC10
NMOS Header	13	Design space exploration	Completed	
	14	Simulation	Completed	
	15	Layout	Completed	
	16	Chip testing	Completed	KAC5, KAC10
Field Programmable Core Array	17	Design space exploration	Dec-2012	KAC11
	18	Synthesis, place and route, noise anlysis	May-2013	KAC12
Energy efficient design tool	19	Add support for conventional (clock gating, power gating, multi-VDD)	Dec-2012	
	20	Add support for proposed (PDVS, resistive power grid, NMOS header)	Dec-2012	
	21	Add support for Encounter Power System	Feb-2013	KAC13

Publications

[KAC1] Shakhsheer, Y., S. Khanna, **K. Craig**, S. Arrabi, J. Lach, and B. H. Calhoun, "A 90nm Data Flow Processor Demonstrating Fine Grained DVS for Energy Efficient Operation from 0.25V to 1.2V", *Custom I ntegrated Circuits Conference*, San Jose, 09/2011.

[KAC2] Calhoun, B. H., S. Arrabi, S. Khanna, Y. Shakhsheer, **K. Craig**, J. Ryan, and J. Lach, "REESES: Rapid Efficient Energy Scalable ElectronicS", *GOMAC Tech*, 03/2010.

[KAC3] Khanna, S., **K. Craig**, Y. Shakhsheer, S. Arrabi, J. Lach, and B. Calhoun, "Stepped Supply Voltage Switching for Energy Constrained Systems", *ISQED*, 2011.

[KAC4] **Craig, K.**, Y. Shakhsheer, S. Khanna, S. Arrabi, J. Lach, B. H. Calhoun, and S. Kosonocky, "A Programmable Resistive Power Grid for Post-Fabrication Flexibility and Energy Tradeoffs", *International Symposium on Low Power Electronics and Design*, 2012.

[KAC5] **Craig, K.**, Y. Shakhsheer, and B. H. Calhoun, "Optimal Power Switch Design for Dynamic Voltage Scaling from High Performance to Subthreshold Operation", *International Symposium on Low Power Electronics and Design*, 2012.

[KAC6] Calhoun, B. H., Y. Zhang, S. Khanna, **K. Craig**, Y. Shakhsheer, J. Lach, "A Sub-Threshold FPGA: Energy-Efficient Reconfigurable Logic."*GOMACTech*. March 2011.

Anticipated publications

[KAC7] **PDVS JSSC journal paper**

[KAC8] BSN revision2 paper.

[KAC9] **PDVS noise analysis**

[KAC10] **Subthreshold power grid design methodologies**

[KAC11] Variable width SIMD paper

[KAC12] **Reconfigurable architecture noise analysis**

[KAC13] **Energy efficient tool paper**

References

[1] M. Putic et al., "Panoptic DVS: A Fine-Grained Dynamic Voltage Scaling Framework for Energy Scalable CMOS Design," *ICCD*, pp.491-497, 2009.

[2] Shakhsheer, Y., S. Khanna, K. Craig, S. Arrabi, J. Lach, and B. H. Calhoun, "A 90nm Data Flow Processor Demonstrating Fine Grained DVS for Energy Efficient Operation from 0.25V to 1.2V", *Custom Integrated Circuits Conference*, San Jose, 09/2011

[3] Craig, K., Y. Shakhsheer, S. Khanna, S. Arrabi, J. Lach, B. H. Calhoun, and S. Kosonocky, "A Programmable Resistive Power Grid for Post-Fabrication Flexibility and Energy Tradeoffs", *International Symposium on Low Power Electronics and Design*, 2012.

[4] Wooters, S., et al., "An Energy-Efficient Subthreshold Level Converter in 130-nm CMOS," *IEEE Trans. Circuits Syst. II* , vol.57, no.4, pp.290-294, 2010.

[5] Craig, K., Y. Shakhsheer, and B. H. Calhoun, "Optimal Power Switch Design for Dynamic Voltage Scaling from High Performance to Subthreshold Operation", *International Symposium on Low Power Electronics and Design*, 2012.

[6] Burd, T., Pering, T., Stratakos, A., Brodersen, R., "A dynamic voltage scaled microprocessor system," *ISSCC*, pp.294-295, 2000.

[7] Jotwani, R., et al., "An x86-64 Core Implemented in 32nm SOI CMOS," *ISSCC*, pp. 106-107, 2010.

Acknowledgements

- Committee Members
  - Mircea Stan, Chair
  - Benton H. Calhoun, Adviser
  - John Lach, Member
  - Joanne Dugan, Member
  - Kevin Skadron, Member
- Fellow Students
  - Yousef, Yanqing, Alicia, Aatmesh, Jim, Craig, Seyi, Peter, Sean, Patricia, Arijit, Divya, Yu, He, Saad, Jeff, Chu, Roman, Sudhanhsu, Sayta, Randy, Joe, Jaijing
- AMD Research
  - Steve Kosonocky